
Replacing the boot media

The boot device stores a primary and secondary set of system files (also called the *boot image*) that the system uses when it boots. The steps you follow to replace a boot device differ depending on whether your system is running ONTAP or 7-Mode.

Before you begin

Before extracting the compressed system files during the boot device installation process, you must make sure that the utility you use is set to suppress any automatic CR/LF conversions in the extracted files.

If you experience any problems during netboot, you can verify the integrity of the files extracted from the package by comparing the `md5sum` values stored in the `MD5SUMS` file.

About this task

- You must replace the failed component with a replacement FRU component you received from your provider.
- The steps for replacing a boot device in a system running 7-Mode are different from a system running ONTAP. You must use the correct procedure for your version of ONTAP 8.x.
- It is important that you apply the commands in these steps on the correct node:
 - The *impaired node* is the node on which you are performing maintenance.
 - The *healthy node* is the HA partner of the impaired node.

Choices

- [Replacing a boot device in a system running ONTAP](#) on page 1
- [Replacing a boot device in a system running in 7-Mode](#) on page 17

Replacing a boot device in a system running ONTAP

Replacing a boot device in ONTAP involves shutting down the node, removing the old boot device, and transferring the system files to the new boot device.

About this task

This procedure is for systems running ONTAP. If your system is running Data ONTAP operating in 7-Mode, you must use that procedure instead.

Steps

1. [Shutting down the target controller](#) on page 2
2. [Opening the system](#) on page 6
3. [Removing the boot device from the controller module](#) on page 8
4. [Installing the boot device](#) on page 9

Shutting down the target controller

You can shut down or take over the target controller by using different procedures, depending on the storage system hardware configuration.

Choices

- [Shutting down a node running ONTAP](#) on page 2
- [Shutting down a node in a two-node MetroCluster configuration running ONTAP](#) on page 5

Shutting down a node running ONTAP

To shut down an impaired node, you must determine the status of the node and, if necessary, take over the node so that the healthy node continues to serve data from the impaired node storage.

Steps

1. If the system is running ONTAP, check the status of the nodes in the cluster.
 - a. Change the privilege level to advanced, entering `y` when prompted to continue: `set -privilege advanced`
The advanced prompt (`*>`) appears.
 - b. Verify the status of the node members in the cluster: `cluster show -epsilon *`

Example

The following example displays information about the health and eligibility of the nodes in the cluster:

Node	Health	Eligibility	Epsilon
node1	true	true	true
node2	true	true	false
node3	true	true	false
node4	true	true	false

4 entries were displayed.

Note: You must not assign epsilon to a node that has to be replaced.

Note: In a cluster with a single HA pair, you must not assign epsilon to either node.

- c. Perform one of the following actions, depending on the result of the command:

If...	Then...
All nodes show <code>true</code> for both health and eligibility, and epsilon is not assigned to the impaired node	<ol style="list-style-type: none">a. Exit advanced mode: <code>set -privilege admin</code>b. Proceed to Step 3.

If...	Then...
All nodes show true for both health and eligibility, and epsilon is assigned to the impaired node	<ol style="list-style-type: none"> a. Remove epsilon from the node: <code>cluster modify -node node1 -epsilon false</code> b. Assign epsilon to a node in the cluster: <code>cluster modify -node node4 -epsilon true</code> c. Exit advanced mode: <code>set -privilege admin</code> d. Go to Step 3.
The impaired node shows false for health and is the epsilon node	<ol style="list-style-type: none"> a. Change the privilege level to advanced: <code>set -privilege advanced</code> b. Remove epsilon from the node: <code>cluster modify -node node1 -epsilon false</code> c. Assign epsilon to a node in the cluster: <code>cluster modify -node node4 -epsilon true</code> d. Exit advanced mode: <code>set -privilege admin</code> e. Proceed to the next step.
The impaired node shows false for health and is not the epsilon node	<ol style="list-style-type: none"> a. Proceed to the next step.
Any node shows false for eligibility	<ol style="list-style-type: none"> a. Resolve any cluster issues as needed. b. Exit advanced mode: <code>set -privilege admin</code>
Any node other than the impaired node shows false for health	<ol style="list-style-type: none"> a. Correct the problems that caused the health issues on the nodes. b. Exit advanced mode: <code>set -privilege admin</code>
<p>2. If the impaired node is part of an HA pair, disable the <code>auto-giveback</code> option from the console of the healthy node: <code>storage failover modify -node local -auto-giveback false</code></p> <p>3. Bring the impaired node to the LOADER prompt:</p>	
If the impaired node is in...	Then...
A stand-alone configuration and is running	Halt the impaired node: <code>system -node halt <i>impaired_node_name</i></code>
A stand-alone configuration and is not running and is not at the LOADER prompt	Resolve any issues that caused the node to quit running, power-cycle it, and then halt the boot process by entering <code>Ctrl-C</code> and responding <code>y</code> to take the node to the LOADER prompt.

If the impaired node is in...	Then...
An HA pair	<p>The impaired node is at the LOADER prompt, it is ready for service. Otherwise, take one of the following actions, as applicable:</p> <ul style="list-style-type: none"> If the impaired node shows the ONTAP prompt, then take over the impaired node from the healthy node and be prepared to interrupt the reboot: <pre>storage failover takeover -ofnode <i>impaired_node_name</i></pre> When prompted to interrupt the reboot, you must press Ctrl-C to go to the LOADER prompt. If the display of the impaired node shows the <code>Waiting for giveback</code> message, then press Ctrl-C and respond y to take the node to the LOADER prompt. If the impaired node does not show either the <code>Waiting for giveback</code> message or an ONTAP prompt, then power-cycle the node. You must contact technical support if the node does not respond to the power cycle.

4. Respond to the applicable wizard:

If LED is...	Then...
Off	<p>NVRAM has no data.</p> <p>Note: You can power down and disconnect the battery.</p>
Flashing	<p>NVRAM is destaging.</p> <p>Note: You must wait for two minutes to complete the destaging operation.</p>
On	<p>NVRAM has data.</p> <p>Note: If you are sure that the data stored in NVRAM is not required, then you can proceed with shut down and replacement. If you have to wipe the NVRAM memory on the impaired controller module, then you must contact technical support for instructions.</p>

5. Shut down the impaired node.

Note: If the system is in an HA pair, the node should be at the LOADER prompt.

The method that you use to shut down the node depends on whether you use remote management through the node's Service Processor (SP), and whether the system is in a dual-chassis configuration or single-chassis configuration.

If the SP is...	Then...
Configured	<p>Log in to the SP of the impaired node, and then turn off the power:</p> <pre>system power off</pre>
Not configured, and the system is in a dual-chassis HA pair in which each controller module is in a separate chassis	Manually shut down the power supplies on the impaired node.
Not configured, and the system is in a single-chassis HA pair in which both controller modules are in the same chassis and share power supplies	At the impaired node prompt, press Ctrl-C and respond y to halt the node.

6. If the system is in a dual-chassis HA pair or stand-alone configuration, turn off the power supplies, and then unplug the power cords of the impaired node from the power source.

Shutting down a node in a two-node MetroCluster configuration running ONTAP

To shut down the impaired node, you must determine the status of the node and, if necessary, switch over the node so that the healthy node continues to serve data from the impaired node storage.

Steps

1. Check the MetroCluster status to determine whether the impaired node has automatically switched over to the healthy node: **metrocluster show**
2. Depending on whether an automatic switchover has occurred, perform one of the following actions:

If the impaired node...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy node: metrocluster switchover
Has not automatically switched over and planned switchover with the metrocluster switchover command fails	<ol style="list-style-type: none">a. Halt the impaired node: system node haltb. Perform a forced switchover operation: metrocluster switchover -forced on disaster true

3. Resynchronize the data aggregates from the surviving cluster: **metrocluster heal -phase aggregates**

Example

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the **metrocluster heal** command with the **-override-vetoes** parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the heal-aggregates operation is completed: **metrocluster operation show**

Example

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
  State: successful
Start Time: 7/25/2014 18:45:55
End Time: 7/25/2014 18:45:56
Errors: -
```

5. Check the state of the aggregates: **storage aggregate show**

Example

```
controller_A_1::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes      RAID Status
-----
...
aggr_b2       227.1GB   227.1GB   0% online    0  mccl-a2    raid_dp, mirrored,
normal...
```

6. Heal the root aggregates: `metrocluster heal -phase root-aggregates`

Example

```
cluster_A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the healing operation is completed on the destination cluster: `metrocluster operation show`

Example

```
cluster_A::> metrocluster operation show
Operation: heal-root-aggregates
State: successful
Start Time: 7/29/2014 20:54:41
End Time: 7/29/2014 20:54:42
Errors: -
```

8. Shut down the impaired node.

The method that you use to shut down the node depends on whether you use remote management through the node's Service Processor (SP), and whether the system is in a dual-chassis configuration or single-chassis configuration.

If the SP is...	Then...
Configured	Log in to the SP of the impaired node, and then turn off the power: system power off
Not configured	Manually shut down the power supplies on the impaired node.

Opening the system

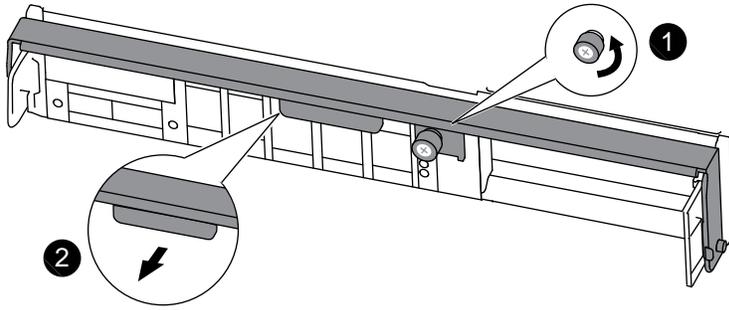
If you want to access components inside the controller module, you must open the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, and keep track of where the cables were connected.

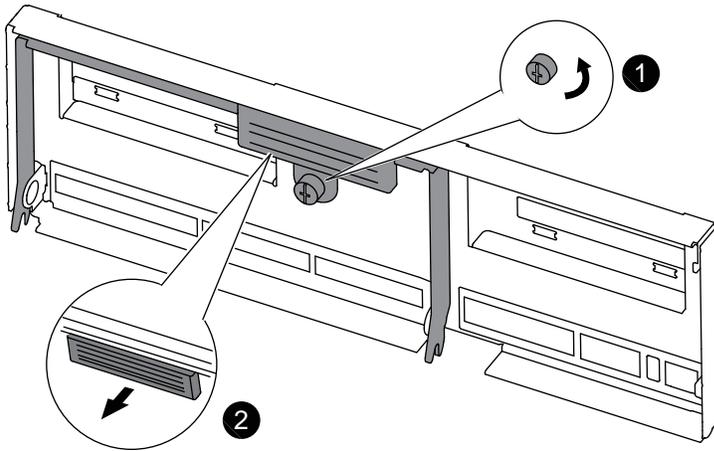
Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
3. Remove the cable management device from the controller module and set it aside.

Note: The 8020 controller module uses cable management arms, while the rest of the 80xx family uses a cable management tray.
4. Loosen the thumbscrew on the cam handle of the controller module.



8020 controller module cam handle

1	Thumbscrew
2	Cam handle



Any other 80xx controller module cam handle

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

If you have...	Then...
An 8020 system	Slide the controller module completely out of the system. Note: Ensure that you support the bottom of the controller module with your free hand and set it aside.
Any other 80xx model	Slide the controller module out of the system until it catches, press the release latch on the left side of the controller module, and then slide the controller module out of the system and set it aside. Note: Ensure that you support the bottom of the controller module with your free hand.

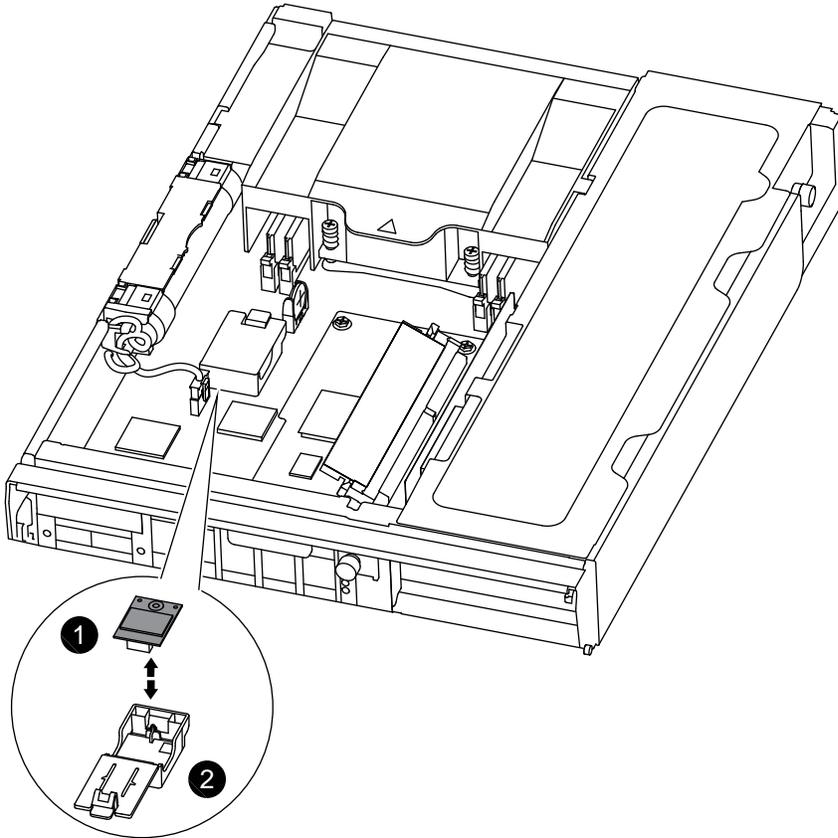
Removing the boot device from the controller module

You must remove the old or failed boot device from the impaired node in your system prior to installing the new or replacement boot device.

Steps

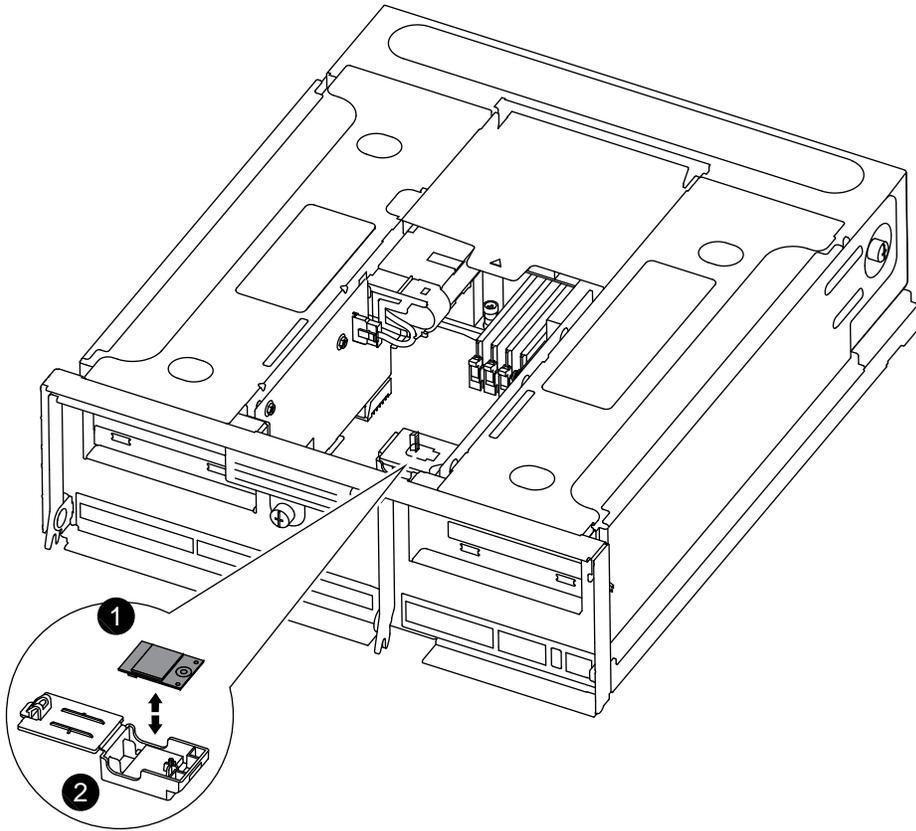
1. If you are not already grounded, properly ground yourself.
2. Remove the controller module and locate the boot device holder.

Use the FRU map on the controller module or the following illustrations to help you locate the boot device enclosure:



8020 controller module

1	Boot device
2	Boot device holder; not removable



Other 80xx controller modules

1	Boot device
2	Boot device holder; not removable

- Open the boot device cover and hold the boot device by its edges at the notches in the boot device housing, gently lift it straight up and out of the housing.

Attention: Always lift the boot device straight up out of the housing. Lifting it out at an angle can bend or break the connector pins in the boot device.

- Set the boot device aside.

Installing the boot device

After you remove the old boot device, you must transfer system files and restore configuration information to your replacement boot device.

Choices

- [Installing the boot device and transferring system files nondisruptively using netboot in systems running Data ONTAP 8.2.1 and later](#) on page 10
- [Installing the boot device and transferring system files disruptively using netboot](#) on page 14

Installing the boot device and transferring system files nondisruptively using netboot in systems running Data ONTAP 8.2.1 and later

You can transfer the system files and restore configuration information to your HA pair nondisruptively by using netboot and the healthy node.

Before you begin

- You must have the blank replacement boot device you received from your provider.
- You must have a network interface to the HA pair healthy node.
- You must have a separate IP address to use to netboot the impaired node.

About this task

The following procedure is written with the assumption that the impaired node has been taken over by the healthy node, and that the impaired node has booted to the LOADER prompt after the takeover has been completed.

Steps

1. Download and extract the `netboot.tgz` file from the NetApp Support Site.

The `netboot.tgz` file is used for performing a netboot of your system. You should download the file contents to a web-accessible directory.

- a. Download the `netboot.tgz` file from the NetApp Support Site to a web-accessible directory.
- b. Switch to the web-accessible directory.
- c. Extract the contents of the `netboot.tgz` file to the target directory `tar -zxvf netboot.tgz`.

Your directory listing should contain the following directory:

```
netboot/
```

Note: If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

2. Download the `image.tgz` file from the NetApp Support Site to the web-accessible directory.

Your directory listing should contain the following file and directory:

```
image.tgz  
netboot/
```

3. If you are not already grounded, properly ground yourself.
4. Turn the controller module so that you can locate the boot device holder.
You can use the FRU map on the controller module to locate the boot device holder.
5. Open the boot device cover, if applicable.
6. Align the boot device with the boot device socket or connector, and then firmly push the boot device straight down into the socket or connector.

Important: Always install the boot device by aligning the front of the boot device squarely over the pins in the socket at the front of the boot device housing. Installing the boot device at an angle or over the rear plastic pin first can bend or damage the pins in the boot device connector.

7. Verify that the boot device is seated squarely and completely in the socket or connector.
If necessary, remove the boot device and reseal it into the socket.
8. Close the boot device cover.
9. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
10. Recable the controller module as needed, and then push the controller module all the way into the chassis.
11. Tighten the thumb screw on the cam handle, reinstall the cable management device, and then tighten the hook and loop strap that binds the cables together.
12. The next step depends on your system configuration:

If your system has...	Then...
One controller module in the chassis	Complete the following substeps: <ol style="list-style-type: none"> a. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power. The system begins to boot and stops at the LOADER prompt. b. Go to the next step.
Two controller modules in the chassis	The node already began rebooting and stops at the LOADER prompt; go to the next step.

13. If you have a switchless cluster, then set the following cluster environment:
 - a. Display the attributes of a switchless cluster from the healthy HA partner:


```
network options switchless-cluster show
```

Note: This is an advanced privilege command.
 - b. Set the switchless cluster environment variable at the LOADER prompt on the impaired node console:


```
setenv bootarg.init.switchless_cluster.enable true
```
14. If you are using Storage Encryption (NSE) check with Self-Encrypting Drives (SED), then contact technical support for assistance with the required bootargs.
15. Set boot environment variables as needed:

If you are running...	Then...
Running Clustered Data ONTAP 8.2.x and later	Set the following boot environment variable at the LOADER prompt on the impaired node console: <pre>setenv bootarg.init.boot_clustered true</pre>
Running ONTAP 8.3 and later on an AFF system	Set the following boot environment variable on the impaired node: <pre>setenv bootarg.init.flash_optimized false</pre>

16. Depending on your network configuration, execute one of the following commands at the LOADER prompt:

If you...	Then...
Have DHCP enabled	Use the following command: <pre>ifconfig e0M -auto</pre>

If you...	Then...
Do not have DHCP enabled	<p>Use the following command:</p> <pre>ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> - dns=<i>dns_addr</i> -domain=<i>dns_domain</i></pre> <p><i>filer_addr</i> is the IP address of the storage system.</p> <p><i>netmask</i> is the network mask of the storage system.</p> <p><i>gateway</i> is the gateway for the storage system.</p> <p><i>dns_addr</i> is the IP address of a name server on your network.</p> <p><i>dns_domain</i> is the Domain Name System (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.</p> <p>Note: To netboot the node when your system is running in Data ONTAP operating in 7-Mode, use an IP address that is not the management IP address for the target. If your system is running ONTAP, you can use the management IP address.</p> <p>Note: Other parameters might be necessary for your interface. For details, use the <code>help ifconfig</code> command at the LOADER prompt.</p>

17. At the LOADER prompt, go to the following web page to install the new software:

```
netboot http://path_to_web-accessible_directory/netboot/kernel
```

18. Select the **Install new software first** option from the displayed menu.

This menu option downloads and installs the new ONTAP image to the boot device.

- You should enter **y** when prompted with the message that this procedure is not supported for nondisruptive upgrade on an HA pair.
- You should enter **y** when warned that this process replaces the existing ONTAP software with new software.
- You should enter the path as follows when prompted for the URL of the *image.tgz* file:

```
http://path_to_the_web-accessible_directory/image.tgz
```

19. Enter **y** when prompted regarding nondisruptive upgrade or replacement of the software.

20. Enter the path to the *image.tgz* file when prompted for the URL of the package.

Example

```
What is the URL for the package?  
http://path_to_web-accessible_directory/image.tgz
```

21. Restore the backup configuration by entering **y** when you see the following screen:

Example

```
*****  
*                Restore Backup Configuration                *  
* This procedure only applies to storage controllers that    *  
* are configured as an HA pair.                             *  
*                                                            *  
* Choose Yes to restore the "varfs" backup configuration    *  
* from the SSH server. Refer to the Boot Device Replacement *  
* guide for more details.                                   *  
* Choose No to skip the backup recovery and return to the  *  
* boot menu.                                               *  
*****
```

```

*****
Do you want to restore the backup configuration
now? {y|n} y

*****
Perform the restore backup procedure on this node's HA partner *
*****

Do you want to restore the backup configuration
now? {y|n} y

Start the restore_backup procedure on this node's HA partner.

```

22. In advanced mode, restore the backup from the HA partner:

a. Set the healthy node to advanced privilege level:

```
set -privilege advanced
```

b. Restore the backup:

```
system node restore-backup -node healthy_node_name -target_address
impaired_node_IP_address
```

c. Return the node to admin level:

```
set -privilege admin
```

23. Restore the configuration from the `varfs.tgz` file, and then reboot the impaired node.

Example

```

The restore configuration has been copied
from the HA partner to this node,
would you like to use this restored copy
now? {y|n} y
Boot device: /dev/da0 found...
Checking /dev/da0s1 file system... success.
Mounting /dev/da0s1 to /cfc card... success.
Checking /cfc card/x86_64/freebsd mount point... success.
Gzip checking /var/home/root/varfs.tgz... success.
Restoring backup configuration...
Backup Configuration from Thu Jun 19 18:01:53 UTC 2014 successfully restored

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y

```

24. After the impaired node has rebooted, perform a giveback from the healthy node:

```
storage failover giveback -ofnode partner_node_name
```

If the giveback is vetoed, you can consider forcing giveback:

```
storage failover giveback -ofnode partner_node_name -require-partner-waiting true -override-
vetoos true
```

The *Clustered Data ONTAP High-Availability Configuration Guide* on the NetApp Support Site contains information about giveback vetoes.

[ONTAP 9 High-Availability Configuration Guide](#)

Installing the boot device and transferring system files disruptively using netboot

You can transfer the system files and restore the configuration information to your HA pair disruptively by performing a netboot to copy the system files to the replacement boot device.

Before you begin

- You must have the blank replacement boot device that you received from your provider.
- You must have access to an HTTP server.
- You must have access to the NetApp Support Site to download the necessary system files for your platform and version of ONTAP software that is running on it.

[NetApp Support](#)

About this task

You can use this disruptive procedure in either a stand-alone system or an HA configuration during scheduled maintenance on the system, or if the controller module is not fully functional.

This procedure is disruptive, even when using it with a system in an HA configuration, because an immediate giveback and re-takeover operation occurs during the procedure. If possible, you should use the nondisruptive procedure for HA configurations for your version of ONTAP software.

Steps

1. Download and extract the `netboot.tgz` file from the NetApp Support Site.

The `netboot.tgz` file is used for performing a netboot of your system. You should download the file contents to a web-accessible directory.

- a. Download the `netboot.tgz` file from the NetApp Support Site to a web-accessible directory.
- b. Switch to the web-accessible directory.
- c. Extract the contents of the `netboot.tgz` file to the target directory `tar -zxvf netboot.tgz`.

Your directory listing should contain the following directory:

```
netboot/
```

Note: If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

2. Download the `image.tgz` file from the NetApp Support Site to the web-accessible directory.

Your directory listing should contain the following file and directory:

```
image.tgz
netboot/
```

3. If you are not already grounded, properly ground yourself.
4. Turn the controller module so that you can locate the boot device holder.
You can use the FRU map on the controller module to locate the boot device holder.
5. Open the boot device cover, if applicable.

- Align the boot device with the boot device socket or connector, and then firmly push the boot device straight down into the socket or connector.

Important: Always install the boot device by aligning the front of the boot device squarely over the pins in the socket at the front of the boot device housing. Installing the boot device at an angle or over the rear plastic pin first can bend or damage the pins in the boot device connector.

- Verify that the boot device is seated squarely and completely in the socket or connector.
If necessary, remove the boot device and reseal it into the socket.
- Close the boot device cover.
- Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
- Recable the controller module as needed, and then push the controller module all the way into the chassis.
- Tighten the thumb screw on the cam handle, reinstall the cable management device, and then tighten the hook and loop strap that binds the cables together.
- The next step depends on your system configuration:

If your system has...	Then...
One controller module in the chassis	Complete the following substeps: <ol style="list-style-type: none"> Reconnect the power cables to the power supplies and to the power sources, and then turn on the power. The system begins to boot and stops at the LOADER prompt. Go to the next step.
Two controller modules in the chassis	The node already began rebooting and stops at the LOADER prompt; go to the next step.

- Display the attributes of a switchless cluster from the healthy HA partner:

```
network options switchless-cluster show
```

- If your node is in a switchless cluster, set the switchless cluster environment variable at the LOADER prompt on the impaired node console:

```
setenv bootarg.init.switchless_cluster.enable true
```

Note: This is an advanced privilege command.

- To set the NetApp Storage Encryption (NSE) boot environment variables, contact technical support.

- Set boot environment variables as needed:

If you are...	Then...
Running ONTAP 8.2.x and later	Set the following boot environment variable at the LOADER prompt on the impaired node console: <pre>setenv bootarg.init.boot_clustered true</pre>
Running ONTAP 8.3.x and later on an AFF system	Set the following boot environment variable on the impaired node: <pre>setenv bootarg.init.flash_optimized true</pre>

- Depending on your network configuration, execute one of the following commands at the LOADER prompt:

If you...	Then...
Have DHCP enabled	Use the following command: ifconfig e0M -auto
Do not have DHCP enabled	Use the following command: ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> - dns=<i>dns_addr</i> -domain=<i>dns_domain</i> <i>filer_addr</i> is the IP address of the storage system. <i>netmask</i> is the network mask of the storage system. <i>gateway</i> is the gateway for the storage system. <i>dns_addr</i> is the IP address of a name server on your network. <i>dns_domain</i> is the Domain Name System (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name. Note: To netboot the node when your system is running in Data ONTAP operating in 7-Mode, use an IP address that is not the management IP address for the target. If your system is running ONTAP, you can use the management IP address. Note: Other parameters might be necessary for your interface. For details, use the <code>help ifconfig</code> command at the LOADER prompt.

18. At the LOADER prompt, boot the system:

```
netboot http://path_to_the_web-accessible_directory/netboot/kernel
```

The system begins to boot, but stops at the Boot menu.

19. Select the **Install new software first** option from the displayed menu.

This menu option downloads and installs the new ONTAP image to the boot device.

- You should enter **y** when prompted with the message that this procedure is not supported for nondisruptive upgrade on an HA pair.
- You should enter **y** when warned that this process replaces the existing ONTAP software with new software.
- You should enter the path as follows when prompted for the URL of the *image.tgz* file:

```
http://path_to_the_web-accessible_directory/image.tgz
```

20. Complete the following substeps:

- Enter **n** to skip the backup recovery when prompted.

```
*****
*                               *
*           Restore Backup Configuration           *
* This procedure only applies to storage controllers that *
* are configured as an HA pair.                    *
*                               *
* Choose Yes to restore the 'varfs' backup configuration *
* from a TFTP server. Refer to the Boot Device Replacement *
* guide for more details.                          *
* Choose No to skip the back up recovery and return to the *
* boot menu.                                       *
*****

Do you want to restore the backup configuration
now? {y|n} n
```

- b. Reboot the node by entering **y** when prompted.

```
The node must be rebooted to start using the newly installed software. Do you want to reboot now? {y/n} y
```

After reboot, if prompted to update firmware and BIOS, enter **y** to accept.

The controller module displays the Boot menu because the boot device was reformatted and the configuration data needs to be restored.

- 21. Select the **Update flash from backup config** option from the displayed menu.

If you are prompted to continue with the update, enter **y**.

- 22. Your next step depends on your system configuration.

If your system is in...	Description
A stand-alone configuration	You can begin using your system after the node reboots to complete restoring the backup configuration.
HA pair	<p>Complete the following substeps after the impaired node is displaying the <code>Waiting for Giveback...</code> message:</p> <ul style="list-style-type: none"> a. Perform a giveback from the healthy node: <pre>storage failover giveback -ofnode replacement_node_name</pre> <p>The impaired node takes back its storage, finishes booting up, and then reboots and is again taken over by the healthy node.</p> <p>Note: If the giveback is vetoed, you can consider overriding the vetoes. ONTAP 9 High-Availability Configuration Guide</p> b. Perform a giveback from the healthy node: <pre>storage failover giveback -ofnode replacement_node_name</pre> <p>The impaired node takes back its storage and finishes booting up with the restored configuration.</p> c. Monitor the progress of the giveback operation by using the <code>storage failover show-giveback</code> command. d. After the giveback operation is complete, confirm that the HA pair is healthy and takeover is possible by using the <code>storage failover show</code> command.

Replacing a boot device in a system running in 7-Mode

Replacing a boot device running in 7-Mode involves shutting down the system, removing the old boot device, and transferring the system files to the new boot device.

About this task

This procedure is for systems running in 7-Mode. If your system is running ONTAP, you cannot use this procedure and must use the procedure for ONTAP instead.

Steps

1. [Shutting down a node running Data ONTAP operating in 7-Mode](#) on page 18
2. [Opening the system](#) on page 19
3. [Removing the boot device from the controller module](#) on page 20

- 4. [Installing the boot device](#) on page 22

Shutting down a node running Data ONTAP operating in 7-Mode

When performing maintenance on a system running Data ONTAP operating in 7-Mode, you must shut down the node. Depending on your system's configuration, you might also need to turn off the power supplies.

About this task

Your system's configuration determines whether you turn off the power supplies after shutting down the node:

- If you have one controller module in the chassis that is either part of an HA pair or in a stand-alone configuration, you must turn off the power supplies in the impaired node chassis.

Shutting down a node in an HA pair

To shut down the node, you must determine the status of the node and, if necessary, take over the node so that the partner continues to serve data from the node's storage.

Steps

1. Check the HA status of the impaired node from either node in the HA pair that is displaying the ONTAP prompt:

```
cf status
```

2. Take the appropriate action based on the takeover status of the node.

If the impaired node...	Then...
Has been taken over by the healthy node and is halted	Go to the next step.
Has not been taken over by the healthy node and is running	Take over the impaired node from the prompt of the healthy node: cf takeover

3. Wait for two minutes after takeover of the impaired node to confirm that the takeover was completed successfully.
4. If the impaired node is not at the LOADER prompt and is showing `waiting for giveback`, press **Ctrl-C** and respond **Y** to halt the node.
5. With the impaired node showing the `waiting for giveback` message or halted, shut it down, depending on your configuration:

If the Service Processor (SP)...	Then...
Is configured	Log in to the SP, and then turn off the power: system power off
Is not configured, and the system is in a dual-chassis HA pair in which each controller is in a separate chassis	Manually shut down the power supplies on the impaired node.
Is not configured, and the system is in a single-chassis HA pair in which both controllers are in the same chassis and share power supplies	At the prompt of the impaired node, press Ctrl-C and respond Y to halt the node.

6. If the nodes are in a dual-chassis HA pair, unplug the impaired node power cords from the power source.

Shutting down a node in a stand-alone configuration

For a node that is not configured with a high-availability (HA) partner, you must perform a clean shutdown (verifying that all data has been written to disk) and disconnect the power supplies.

Steps

1. Shut down the node if it is not already shut down:

```
halt -t 0
```
2. Shut down the power supplies, and then unplug both power cords from the source.
The system is ready for maintenance.

Opening the system

If you want to access components inside the controller module, you must open the system.

Steps

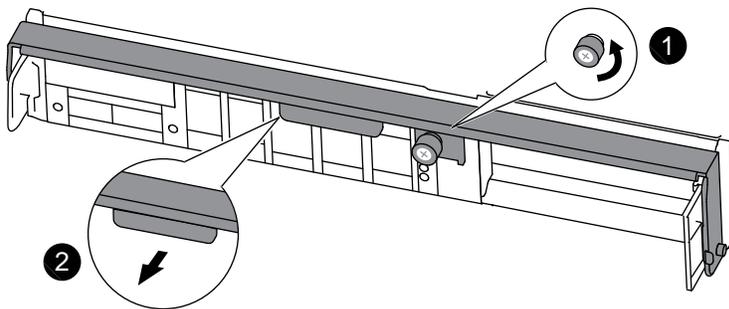
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, and keep track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove the cable management device from the controller module and set it aside.

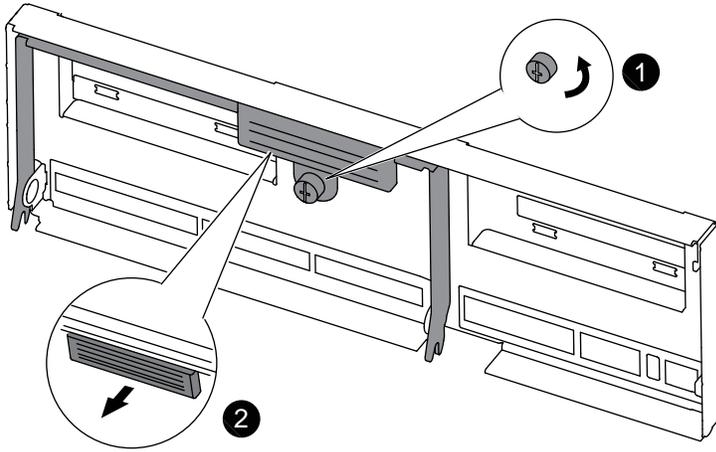
Note: The 8020 controller module uses cable management arms, while the rest of the 80xx family uses a cable management tray.

4. Loosen the thumbscrew on the cam handle of the controller module.



8020 controller module cam handle

1	Thumbscrew
2	Cam handle



Any other 80xx controller module cam handle

1	Thumbscrew
2	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

If you have...	Then...
An 8020 system	Slide the controller module completely out of the system. Note: Ensure that you support the bottom of the controller module with your free hand and set it aside.
Any other 80xx model	Slide the controller module out of the system until it catches, press the release latch on the left side of the controller module, and then slide the controller module out of the system and set it aside. Note: Ensure that you support the bottom of the controller module with your free hand.

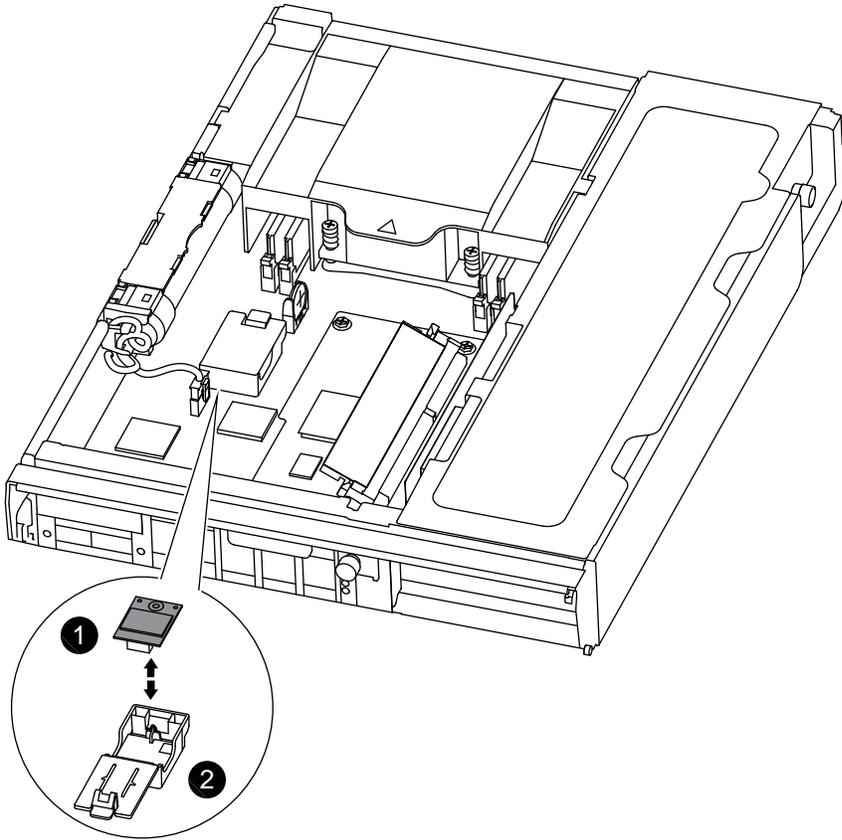
Removing the boot device from the controller module

You must remove the old or failed boot device from the impaired node in your system prior to installing the new or replacement boot device.

Steps

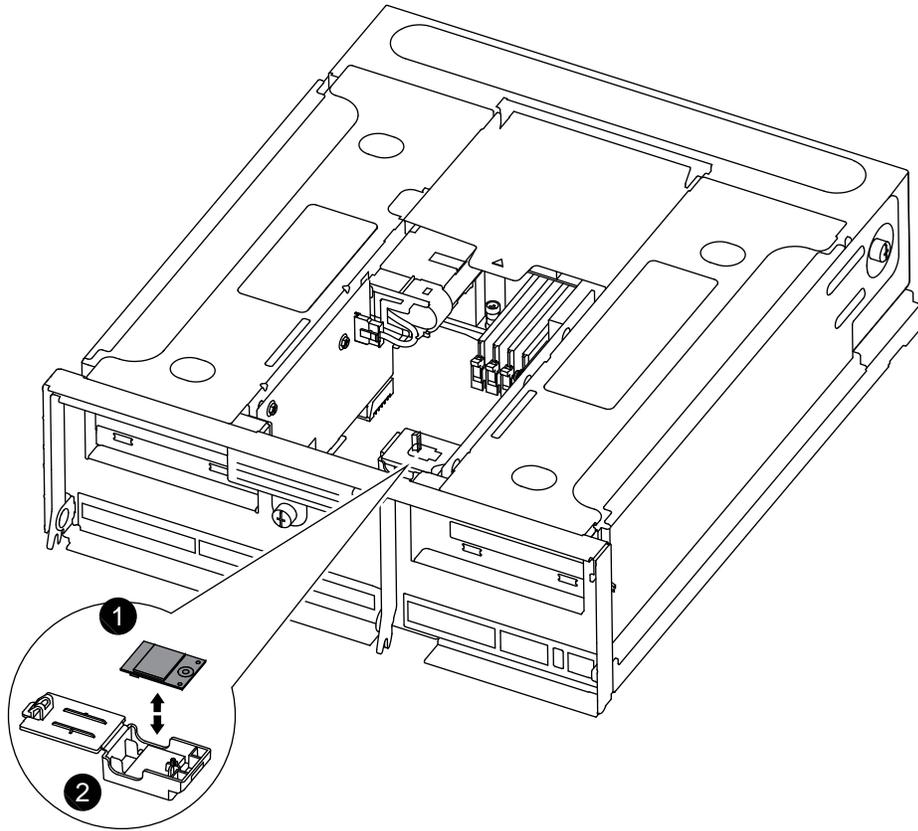
1. If you are not already grounded, properly ground yourself.
2. Remove the controller module and locate the boot device holder.

Use the FRU map on the controller module or the following illustrations to help you locate the boot device enclosure:



8020 controller module

<p>1</p>	<p>Boot device</p>
<p>2</p>	<p>Boot device holder; not removable</p>



Other 80xx controller modules

1	Boot device
2	Boot device holder; not removable

- Open the boot device cover and hold the boot device by its edges at the notches in the boot device housing, gently lift it straight up and out of the housing.

Attention: Always lift the boot device straight up out of the housing. Lifting it out at an angle can bend or break the connector pins in the boot device.

- Set the boot device aside.

Installing the boot device

After you remove the old boot device, you must transfer system files and restore configuration information to your replacement boot device.

Choices

- [Installing the boot device and transferring system files nondisruptively using netboot in systems running ONTAP 8.2.1 and later on page 23](#)
- [Installing the boot device and transferring system files disruptively using netboot on page 26](#)

Installing the boot device and transferring system files nondisruptively using netboot in systems running ONTAP 8.2.1 and later

You can transfer the system files and restore configuration information to your HA pair nondisruptively by using netboot and the healthy node.

Before you begin

- You must have the blank replacement boot device that you received from your provider.
- You must have a network interface to the HA pair healthy node.
- You must have a separate IP address to use to netboot the impaired node.

About this task

The following procedure is written with the assumption that the impaired node has been taken over by the healthy node, and that the impaired node has booted to the LOADER prompt after the takeover has been completed.

Steps

1. Download and extract the `netboot.tgz` file from the NetApp Support Site.

The `netboot.tgz` file is used for performing a netboot of your system. You should download the file contents to a web-accessible directory.

- a. Download the `netboot.tgz` file from the NetApp Support Site to a web-accessible directory.
- b. Switch to the web-accessible directory.
- c. Extract the contents of the `netboot.tgz` file to the target directory `tar -zxvf netboot.tgz`.

Your directory listing should contain the following directory:

```
netboot/
```

Note: If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

2. Download the `image.tgz` file from the NetApp Support Site to the web-accessible directory.

Your directory listing should contain the following file and directory:

```
image.tgz
netboot/
```

3. If you are not already grounded, properly ground yourself.
4. Turn the controller module so that you can locate the boot device holder.
You can use the FRU map on the controller module to locate the boot device holder.
5. Open the boot device cover, if applicable.
6. Align the boot device with the boot device socket or connector, and then firmly push the boot device straight down into the socket or connector.

Important: Always install the boot device by aligning the front of the boot device squarely over the pins in the socket at the front of the boot device housing. Installing the boot device at an angle or over the rear plastic pin first can bend or damage the pins in the boot device connector.

7. Verify that the boot device is seated squarely and completely in the socket or connector.
If necessary, remove the boot device and reseal it into the socket.
8. Close the boot device cover.
9. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
10. Tighten the thumb screw on the cam handle, reinstall the cable management device, and then tighten the hook and loop strap that binds the cables together.

11. The next step depends on your system configuration:

If your system has...	Then...
One controller module in the chassis	Complete the following substeps: <ol style="list-style-type: none"> a. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power. The system begins to boot and stops at the LOADER prompt. b. Go to the next step.
Two controller modules in the chassis	The node already began rebooting and stops at the LOADER prompt; go to the next step.

12. Depending on your network configuration, execute one of the following commands at the LOADER prompt:

If you...	Then...
Have DHCP enabled	Use the following command: ifconfig e0M -auto
Do not have DHCP enabled	Use the following command: ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> - dns=<i>dns_addr</i> -domain=<i>dns_domain</i> <i>filer_addr</i> is the IP address of the storage system. <i>netmask</i> is the network mask of the storage system. <i>gateway</i> is the gateway for the storage system. <i>dns_addr</i> is the IP address of a name server on your network. <i>dns_domain</i> is the Domain Name System (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name. Note: To netboot the node when your system is running in Data ONTAP operating in 7-Mode, use an IP address that is not the management IP address for the target. If your system is running ONTAP, you can use the management IP address. Note: Other parameters might be necessary for your interface. For details, use the <code>help ifconfig</code> command at the LOADER prompt.

13. At the LOADER prompt, netboot the kernel:

```
netboot http://path_to_web-accessible_directory/netboot/kernel
```

14. Select the **Install new software first** option from the displayed menu.

This menu option downloads and installs the new ONTAP image to the boot device.

- You should enter **y** when prompted with the message that this procedure is not supported for nondisruptive upgrade on an HA pair.

- You should enter **y** when warned that this process replaces the existing ONTAP software with new software.
- You should enter the path as follows when prompted for the URL of the *image.tgz* file:

```
http://path_to_the_web-accessible_directory/image.tgz
```

15. Enter **y** when prompted regarding nondisruptive upgrade or replacement of the software.
16. Enter the path to the *image.tgz* file when prompted for the URL of the package.

Example

```
What is the URL for the package?
http://path_to_web-accessible_directory/image.tgz
```

17. Restore the backup configuration on the impaired node by entering **y**, and then entering the IP address of the server containing the *varfs.tgz* file when you see the following screen:

```
*****
*                               *
*       Restore Backup Configuration       *
* This procedure only applies to storage controllers that *
* are configured as an HA pair.           *
*                               *
* Choose Yes to restore the "varfs" backup configuration *
* from the SSH server. Refer to the Boot Device Replacement *
* guide for more details.                 *
* Choose No to skip the backup and recovery and return to the *
* boot menu.                               *
*****
Do you want to restore the backup configuration
now? {y|n} y

Start the restore_backup procedure on this node's HA partner.
```

18. Restore the backup configuration from the HA partner:
restore_backup target_node.netboot_ip.address
19. Restore the configuration from the *varfs.tgz* file and reboot the impaired node, as shown in the following example:

Example

```
The restore configuration has been copied
from the HA partner to this node,
would you like to use this restored copy
now? {y|n} y
Boot device: /dev/da0 found...
Checking /dev/da0s1 file system... success.
Mounting /dev/da0s1 to /cfcard... success.
Checking /cfcard/x86_64/freebsd mount point... success.
Gzip checking /var/home/root/varfs.tgz... success.
Restoring backup configuration...
Backup Configuration from Thu Jun 19 18:01:53 UTC 2014 successfully restored

The node must be rebooted to start using the newly installed software. Do you
want to reboot now? {y|n} y
```

20. After the impaired node has rebooted and is displaying the *Waiting for giveback* message, complete the procedure by giving back the node from the healthy node console:

cf giveback

Installing the boot device and transferring system files disruptively using netboot

You can transfer the system files and restore configuration information to your system by using netboot to copy the system files to the replacement boot device.

Before you begin

- You must have the blank replacement boot device that you received from your provider.
- You must have access to an HTTP server.
- You must have access to the NetApp Support Site at mysupport.netapp.com.
This enables you to download the necessary system files for your platform and version of ONTAP running on it.

About this task

You can use this disruptive procedure in either a stand-alone system or an HA configuration during scheduled maintenance on the system or if the healthy node is not fully functional.

This procedure is disruptive even when using it with a system in an HA configuration, because an immediate giveback and re-takeover occur during the procedure. If possible, you should use the nondisruptive procedure for HA configurations for your version of Data ONTAP software operating in 7-Mode.

Steps

1. Download and extract the `netboot.tgz` file from the NetApp Support Site.

The `netboot.tgz` file is used for performing a netboot of your system. You should download the file contents to a web-accessible directory.

- a. Download the `netboot.tgz` file from the NetApp Support Site to a web-accessible directory.
- b. Switch to the web-accessible directory.
- c. Extract the contents of the `netboot.tgz` file to the target directory `tar -zxvf netboot.tgz`.

Your directory listing should contain the following directory:

```
netboot/
```

Note: If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

2. Download the `image.tgz` file from the NetApp Support Site to the web-accessible directory.

Your directory listing should contain the following file and directory:

```
image.tgz
netboot/
```

3. If you are not already grounded, properly ground yourself.
4. Turn the controller module so that you can locate the boot device holder.
You can use the FRU map on the controller module to locate the boot device holder.
5. Open the boot device cover, if applicable.

6. Align the boot device with the boot device socket or connector, and then firmly push the boot device straight down into the socket or connector.

Important: Always install the boot device by aligning the front of the boot device squarely over the pins in the socket at the front of the boot device housing. Installing the boot device at an angle or over the rear plastic pin first can bend or damage the pins in the boot device connector.

7. Verify that the boot device is seated squarely and completely in the socket or connector.
If necessary, remove the boot device and reseal it into the socket.
8. Close the boot device cover.
9. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
10. Recable the controller module as needed, and then push the controller module all the way into the chassis.
11. Tighten the thumb screw on the cam handle, reinstall the cable management device, and then tighten the hook and loop strap that binds the cables together.
12. The next step depends on your system configuration:

If your system has...	Then...
One controller module in the chassis	Complete the following substeps: <ol style="list-style-type: none"> a. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power. The system begins to boot and stops at the LOADER prompt. b. Go to the next step.
Two controller modules in the chassis	The node already began rebooting and stops at the LOADER prompt; go to the next step.

13. Depending on your network configuration, execute one of the following commands at the LOADER prompt:

If you...	Then...
Have DHCP enabled	Use the following command: ifconfig e0M -auto
Do not have DHCP enabled	Use the following command: ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> - dns=<i>dns_addr</i> -domain=<i>dns_domain</i> <i>filer_addr</i> is the IP address of the storage system. <i>netmask</i> is the network mask of the storage system. <i>gateway</i> is the gateway for the storage system. <i>dns_addr</i> is the IP address of a name server on your network. <i>dns_domain</i> is the Domain Name System (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name. Note: To netboot the node when your system is running in Data ONTAP operating in 7-Mode, use an IP address that is not the management IP address for the target. If your system is running ONTAP, you can use the management IP address. Note: Other parameters might be necessary for your interface. For details, use the <code>help ifconfig</code> command at the LOADER prompt.

14. At the LOADER prompt, boot the system:

```
netboot http://path_to_the_web-accessible_directory/netboot/kernel
```

The system begins to boot, but stops at the Boot menu.

15. Select the **Install new software first** option from the displayed menu.

This menu option downloads and installs the new ONTAP image to the boot device.

- You should enter **y** when prompted with the message that this procedure is not supported for nondisruptive upgrade on an HA pair.
- You should enter **y** when warned that this process replaces the existing ONTAP software with new software.
- You should enter the path as follows when prompted for the URL of the *image.tgz* file:

```
http://path_to_the_web-accessible_directory/image.tgz
```

16. Complete the following substeps:

a. Enter **n** to skip the backup recovery when you see the following prompt:

```
*****
*           Restore Backup Configuration           *
* This procedure only applies to storage controllers that *
* are configured as an HA pair.                    *
*                                                    *
* Choose Yes to restore the 'varfs' backup configuration *
* from a TFTP server. Refer to the Boot Device Replacement *
* guide for more details.                          *
* Choose No to skip the back up recovery and return to the *
* boot menu.                                        *
*****

Do you want to restore the backup configuration
now? {y|n} n
```

b. Reboot the node by entering **y** when you see the following prompt:

```
The node must be rebooted to start using the newly installed
software. Do you want to reboot now? {y/n} y
```

The controller module reboots, but stops at the Boot menu because the boot device was reformatted and the configuration data needs to be restored.

17. Select the **Update flash from backup config** option from the displayed menu.

18. Enter **y** when prompted.

The system automatically reboots when the update flash (syncflash) process is complete.

```
This will replace all flash-based configuration with the last backup to
disk. Are you sure you want to continue?: y
```

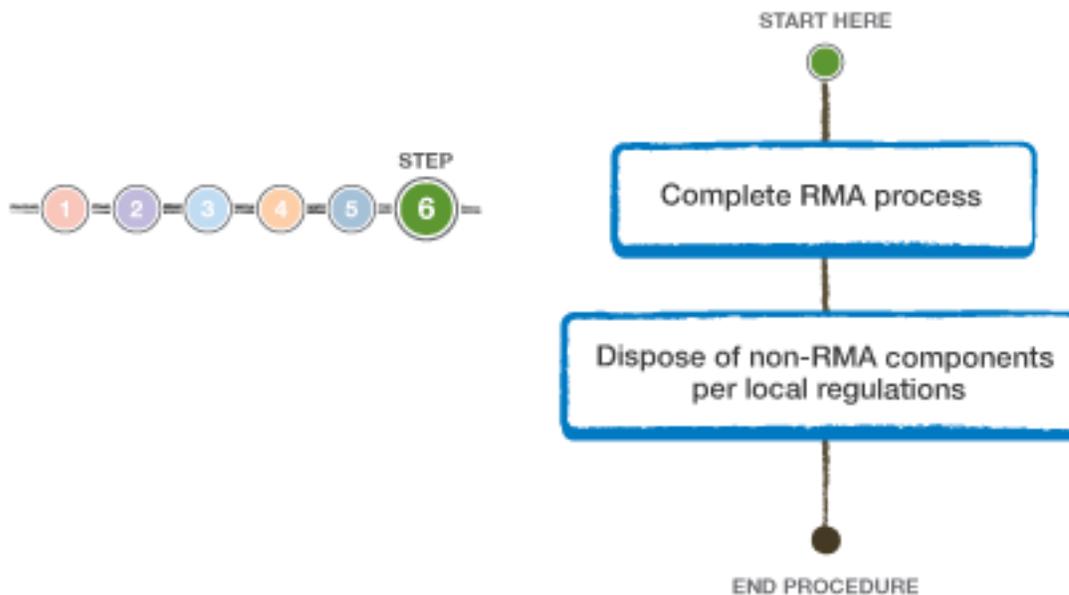
19. Your next step depends on your system configuration.

If your system is...	Description
A stand-alone configuration	You can begin using your system after the node reboots to complete restoring the backup configuration.

If your system is...	Description
An HA pair	<p>After the impaired node displays the <code>Waiting for Giveback...</code> message, perform a giveback from the healthy node:</p> <pre>cf giveback</pre> <p>Note: The impaired node takes back its storage, finishes booting up, and then reboots and is again taken over by the healthy node.</p>

Completing the replacement process

After you replace the part, you can return the failed part to NetApp, as described in the RMA instructions shipped with the kit. Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.



Copyright

Copyright © 2019 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR

PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277